

Human-interaction challenges in UAV-based autonomous surveillance

Michael Freed

freed@arc.nasa.gov
NASA Ames Research Center
Institute for Human and Machine Cognition

Robert Harris

rharris@arc.nasa.gov
NASA Ames Research Center
QSS Inc.

Michael G. Shafte

mshafto@arc.nasa.gov
NASA Ames Research Center

Abstract

Autonomous UAVs provide a platform for intelligent surveillance in application domains ranging from security and military operations to scientific information gathering and land management. Surveillance tasks are often long duration, requiring that any approach be adaptive to changes in the environment or user needs. We describe a decision-theoretic model of surveillance, appropriate for use on our autonomous helicopter, that provides a basis for optimizing the value of information returned by the UAV. From this approach arise a range of challenges in making this framework practical for use by human operators lacking specialized knowledge of autonomy and mathematics. This paper describes our platform and approach, then describes human-interaction challenges arising from this approach that we have identified and begun to address.

UAV-based Surveillance

One of the earliest applications of powered air vehicles was to gather information about conditions on the ground, exploiting the relatively high speed and broad view offered by these machines to, e.g., provide guidance to World War I artillery units and track enemy movements. Similar applications quickly emerged in other areas such as security, land management and scientific research. Unmanned aerial vehicles (UAVs) have the potential to dramatically increase the availability and usefulness of aircraft as information-gathering platforms. As UAV technologies improve and the number of such vehicles increases, costs will come to reflect economy of scale and decreased weight and complexity made possible by not having to support human occupants. Reduced cost should, in turn, make UAVs available to a wider and less specialized set of users and for increasingly diverse purposes. This presents a challenge: how best to accommodate increasingly diverse missions and users of UAV-based observation platforms.

UAVs in operational use have been employed mainly for reconnaissance and surveillance (DoD 2002).

Reconnaissance tasks are intended to acquire information about a particular area at a particular time, generally to inform a specific impending decision. We generalize the normal usage slightly in our work to include any once-only observation task. By this definition, a UAV helicopter assessing the suitability of a potential landing spot would be performing reconnaissance, as it would when handling a human user's request to check a building entrance for intruders. *Surveillance* tasks differ from reconnaissance tasks in that they involve repeated or continuous observation intended to maintain awareness of some entity or geographical area. Because surveillance generally takes place over a lengthy period, it is particularly appropriate to carry out these missions with autonomous vehicles rather than, as is now the norm, with humans as remote pilots.

However, performing surveillance autonomously is particularly challenging for several reasons. One is that, for surveillance of multiple, spatially-separated targets, the core autonomy problem of deciding where to observe next is inherently a difficult scheduling problem. Choosing to observe one site has a cascading effect on both the time-cost and desirability of all subsequent observation tasks. This coupling of current and future choices suggests considering each observation to be part of a schedule whose utility can be compared to possible alternative schedules. However, unlike typical AI scheduling problems, the surveillance problem requires allowing for the possibility that some sites should be visited more often than others due to differences in their importance and in the rates at which observed information becomes obsolete. Given an overall goal of maximizing the value of information returned to the user, a surveillance scheduler should, in many cases, omit visits to some (possibly most) sites entirely in order to observe the most important ones at a higher rate. Surveillance scheduling thus combines task ordering with task selection, a combination notorious for increasing the computational complexity of any solution.

Many factors can have a dramatic effect on the effectiveness of surveillance activities and must therefore be taken into account during the scheduling process. For instance, evaluating the desirability of a candidate site will typically require estimating how long it will take to get to the site and take any needed measurements. However, even assuming that the site is at a known, fixed location,

UAV-relevant variables such as weather, initial velocity and required turn angle from the last visited site can all have a large impact on traverse time. Other factors affect what kind of approach will be effective in comparing alternatives. For instance, an algorithm that works well when the number of surveillance sites is small (say, 5) may not work well when the number of sites is an order magnitude larger. Similarly, algorithms that attempt to take advantage of problem structure – e.g. spatial “clumping” of sites or non-uniform distribution of event probabilities across these sites – will not be effective in problem instances lacking these structural features. Current work in autonomous surveillance falls far short in dealing with the range of qualitatively distinct surveillance scheduling problems that might be encountered in real operational contexts.

The Autonomous Rotorcraft Project

Our work on surveillance scheduling is one element of the Autonomous Rotorcraft Project (Whalley et al. 2003), an Army/NASA effort to develop high-level autonomy for airborne observation missions of interest to both organizations. The vehicle selected for the project, a Yamaha RMAX helicopter (see Figure 1), can fly at low speed or remain in a hover for approximately 60 minutes while carrying a 65lb. payload. Having completed an initial round of autonomous flight tests using newly developed flight control software, we expect that the vehicle will soon be capable of prolonged autonomous flight in an agile, aerodynamic flight mode. This will extend its effective speed, range and maximum flight duration and thus, we believe, make it practical as an Intelligent Surveillance and Reconnaissance (ISR) platform and as a demonstration of how ISR capabilities might be implemented on other platforms such as fixed wing aircraft and motorized blimps.



Figure 1. Project platform: Yamaha RMAX

High-level autonomous control is provided by Apex, a reactive, procedure-based planner/scheduler used for mission-level task execution, navigation, response to health and safety contingencies and interaction with human users. Surveillance scheduling in a realistically dynamic mission context – i.e. where flight conditions and user needs can change often and unexpectedly – is seen as a special case of the problem of multitask management, a central Apex capability and research focus (Freed 1998, 2000). Though the approach it incorporates has proven effective for some relatively complex tasks (John et al. 2001; Freed et al. 2003), the surveillance problem has proven much more demanding. Our work on autonomous surveillance has therefore focused on identifying or designing candidate surveillance scheduling approaches, characterizing their strengths and weaknesses, and incorporating the most effective approaches into Apex. We allow for the possibility that human-directed surveillance – i.e. where a human operator decides where to observe next – will prove more effective than any known algorithm in some circumstances.

A key part of our effort is a framework for evaluating surveillance decision performance in a wide range of mission scenarios. Like Massios, Dorst and Voorback (2001), we take a decision-theoretic approach to defining the surveillance problem. The value of making an observation at a particular time and place, then returning that information to a user, depends on the kinds of events that might be observed and the value the user places on knowing about them. As the value of information often depends on when the user receives it (e.g. it is better to be informed of a break-in as it is beginning than long after the thief has escaped), surveillance decisions should take into account temporal characteristics of the task environment such as the likelihood of an interesting event occurring over a given interval and the change over time in the value of observing that event after it occurs. Our approach treats observations as boundaries on time-intervals in which the user has been ignorant of events occurring at a given site (target). The expected cost of ignorance (ECI) for a given target over a given interval is:

$$ECI_{\tau}(t_1, t_2) = \int_{t=t_1}^{t_2} p(t) \cdot cost(t_2 - t) dt$$

where (t_1, t_2) is the interval between observations measured from some mission start time t_0 , $p(t)$ is probability density function for the occurrence of some cost-imposing event E (e.g. a break-in) and $cost(d)$ is a function describing the expected cost imposed by E as a function of the time from occurrence to intervention. The expected cost of ignorance is thus the sum, for all points in the interval, of the probability of the event occurring at that point times the expected cost if it occurs at that point. With cost and probability functions appropriate to model events of type

E, the total cost of a given surveillance schedule is the sum of ECIs for all inter-observation intervals for all targets. The value of observations resulting from following the schedule is the worst case schedule cost (no observations at all) minus the cost of the selected schedule. The goal of a surveillance algorithm is to maximize that value. See Freed et al. (2004) for more detail on this approach.

Human-Interaction Challenges

As described below, we have identified a range of human-interaction issues arising in from our decision-theoretic approach to autonomous surveillance and have made some progress addressing the first of these issues.

Role of humans in deciding surveillance pattern

Surveillance is often a lengthy, repetitive and largely uneventful process that strains human vigilance and morale. This makes it an excellent application for autonomy. However, it is difficult to formulate algorithms that perform well at full range of possible surveillance missions. This suggests that it may sometimes be desirable to have human operators select and order surveillance observations. To better understand human strength and weaknesses in making these decisions, and thus to become better able to allocate responsibilities between an autonomous controller and human operator, we compared human performance against that of a surveillance algorithm using a set of 243 mission scenarios. These scenarios, intended to represent a significant portion of the diversity of surveillance missions in realistic NASA and Army contexts, differ in 5 dimensions (independent variables), each with 3 possible values:

Number of observation targets: 4, 8 or 16
Spatial scale: .002, .02, .2 of vehicle range
Spatial distribution: uniform, globular, 2-cluster
Cost max distribution: fixed, uniform, 2-cluster
Cost rate distribution: fixed, uniform, 2-cluster

As a candidate autonomous surveillance decision-method, the study used a variant of the 2-opt algorithm (Reinelt 1994) for the Traveling Salesman Problem modified to generate repeating sequences and to use a history-sensitive UAV kinematics model to compute travel time between targets. Schedules for each scenario generated by this method were evaluated using the ECI approach outlined above.

For this initial study, our assessment of human surveillance performance was based on data from a single subject. Each scenario was depicted as a map (see Figure 2) with observation targets colored-coded to indicate cost-rate (urgency) and shape-coded to indicate maximum cost (importance). The start/end point (home) was depicted as a distinctive icon and spatial scale as a dotted-line circle,

centered on the home point, whose radius represented a fixed proportion of the vehicle's range. The subject used a mouse to select and modify a route. The amount of time taken to select each route was displayed, though no time limit was enforced. The subject was very familiar with the surveillance task but was not given training on effective strategy or provided with any decision aids. Comparison between human and algorithm performance (see Table 1) showed a number of trends including the following:

- The human subject performed poorly in scenarios where surveillance targets were nearby one another (small spatial scale). In these cases, optimal paths are often complex and unintuitive due to the relatively large effect of vehicle kinematics on travel time.
- The subject did relatively well when the number of targets is large (16), but poorly with few targets (4) where the importance of small differences in optimality gets magnified by repetition
- The subject performed well where there was a lot of structure to reason about, especially including spatial structure (sets of targets "clumped" together), but also structure in the distribution of urgency and importance; in contrast, the subject did poorly when targets were spaced uniformly (no clumps) and with uniform distributions of other features

Future studies including a larger number of human subjects provided with training and decision aids will enable a more reliable and detailed assessment of relative strengths and weaknesses.

Handling reconnaissance requests

The problem of how to choose where to observe next has been the focus of our efforts so far. However, other issues relating to the need to interact with human users clearly must be addressed if the autonomous surveillance approach we are developing is to become practical for widespread use. First, most practical applications will require allowing users to make reconnaissance requests that must be executed against the "background" surveillance task. Extending our surveillance algorithms (we are combining several) to insert reconnaissance actions into the current schedule only addresses part of the problem since there are many ways to accommodate such a request. A requestor might mean "drop everything and go perform this observation" but might also be making a more nuanced request such as "make this observation at the most convenient time in the current schedule" or "optimize scheduling of this reconnaissance observation using the same framework used to schedule surveillance observations." To handle these different kinds of requests, the system needs to engage in some kind of dialogue with the user to understand what degrees of freedom exist in

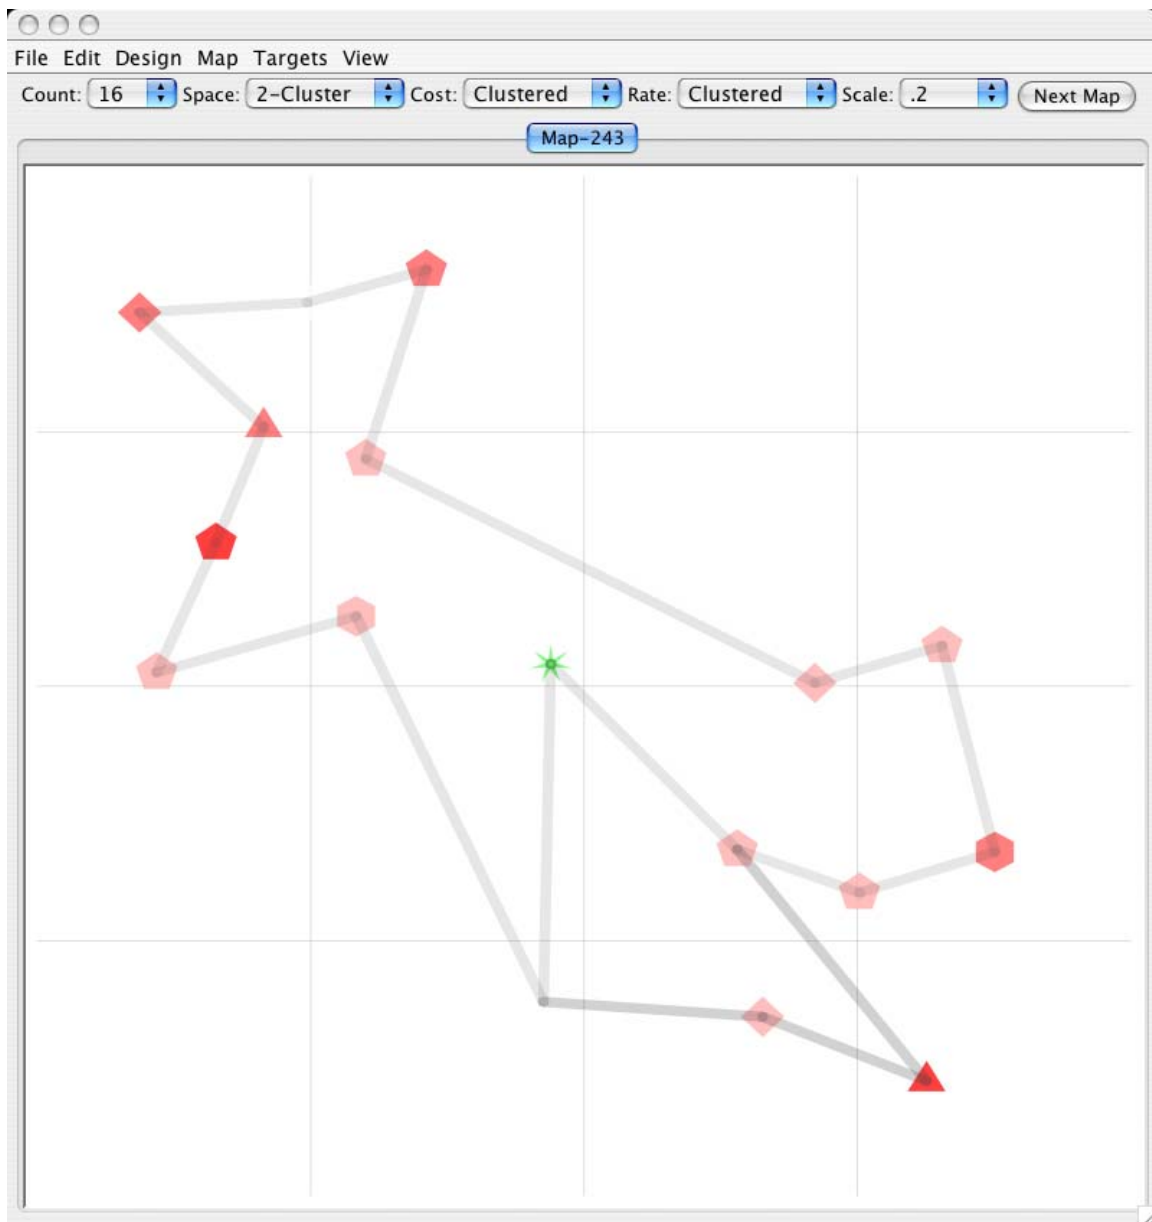


Figure 2. One of 243 scenarios used to evaluate human/algorithm performance

scheduling the proposed action and to insure that the user understands and accepts the observation delays that will result from carrying it out.

Acquiring model data from the user

Our approach to surveillance scheduling relies on quantitative, temporally structured models of the user and environment. For instance, an autonomous surveillance agent might be given a set of observation goals that include checking for fires in a particular region of forest. Quantifying the expected cost of remaining ignorant of the state of that site requires a model of how a fire there might spread and how costly the user considers a fire that has

spread for a given interval. To characterize the mounting cost of a forest fire, the user might realize that fires tend to start slowly and then spread rapidly, but eventually slow down as they consume the fuel in a region and approach natural boundaries. A mathematician might choose to model this as an appropriately parameterized sigmoid function.

An autonomous vehicle in daily use for fire surveillance is likely to be attended by people who know a lot about forests and forest fires, but are not mathematicians. Nonetheless, they may need to construct such models and modify them periodically. For instance, seasonal changes in temperature and humidity, new construction, and manual thinning of the tree canopy or

Pct. Adv.			N Spatial								
			4			8			16		
Scale	Rate	Cost	2-Cluster	Globular	Uniform	2-Cluster	Globular	Uniform	2-Cluster	Globular	Uniform
0.002	Fixed	Fixed	0	0	0	0	14	0	6	31	-3
		Clustered	0	0	0	0	16	5	7	29	7
		Uniform	0	0	-1	42	24	5	2	10	10
	Clustered	Fixed	0	0	0	0	0	11	0	0	11
		Clustered	0	0	0	0	0	0	0	0	1
		Uniform	0	0	0	0	0	0	0	0	1
	Uniform	Fixed	143	47	0	130	19	47	10	93	47
		Clustered	22	28	23	0	12	6	3	8	0
		Uniform	61	20	-2	0	18	7	2	5	9
0.02	Fixed	Fixed	0	0	0	1	1	0	0	1	0
		Clustered	0	0	0	1	1	0	0	1	0
		Uniform	0	0	0	1	1	0	1	2	0
	Clustered	Fixed	0	0	0	1	0	0	1	6	0
		Clustered	0	0	0	1	0	0	0	9	1
		Uniform	0	0	0	1	16	0	0	6	0
	Uniform	Fixed	0	0	0	1	0	0	2	2	0
		Clustered	0	0	0	1	0	0	2	2	1
		Uniform	0	0	0	1	0	0	0	2	1
0.2	Fixed	Fixed	1	0	0	0	-1	0	-2	-22	4
		Clustered	1	0	0	-1	5	-10	-6	-15	3
		Uniform	0	0	-1	-2	4	-8	-5	-19	4
	Clustered	Fixed	0	0	0	8	0	14	9	-1	10
		Clustered	0	0	0	9	-1	-11	11	-14	5
		Uniform	0	0	-1	16	8	3	5	-10	4
	Uniform	Fixed	3	0	23	14	-7	22	2	-7	10
		Clustered	16	0	23	9	-3	-12	9	-16	8
		Uniform	23	0	31	15	4	2	-3	-26	5

Table 1. Percentage difference in performance between 2-Opt and human-directed surveillance. Positive values indicate advantage for 2-Opt. Boldface indicates a difference $\geq 10\%$.

ground can all affect the fire spreading characteristics of a region of forest. One possibility is to create a specialized interface for forest fire surveillance that facilitates changes to the model, but such specialization has drawbacks. First, it implies the need for expensive, highly specialized expertise to formulate an initial model and to provide a customized interface that supports users in adapting the model. Second, it reduces the flexibility of the platform for any given user, essentially requiring that they use it just as originally intended. A more general approach would allow users to adapt not only to changes in the environment, but also to changing needs. For example, forest management personnel may want to use a UAV mainly for fire surveillance, but also for additional jobs such as gathering scientific data for a long-term wildlife study or periodically checking vegetation conditions at remote sites. Using a general approach makes it possible to take on qualitatively new surveillance goals without making fundamental changes to the system.

Decision-theoretic characterizations of user information needs and environment attributes provide a basis for optimal scheduling of surveillance actions, but are not necessarily intuitive. Over an extended period of operation, the need to support users in modifying the models used for reconnaissance scheduling takes on added importance. Personnel changes may cause loss of

knowledge about the rationale underlying existing models. The capabilities and operational characteristics of the vehicle may change due to upgrades (or wear and tear). And, generally, the more time passes, the more reality will “drift” from what was captured in the original models. Addressing this problem – i.e. bridging the gap between user conceptualization of a task domain and mathematically precise but unintuitive representations of that domain is thus a key issue in providing intelligent surveillance capabilities for long duration missions in dynamic environments.

Coordination and integration

A third human-interaction issue is the need to integrate the autonomous surveillance agent into a larger system, coordinating its activities with those of other agents whose actions might prove helpful, harmful or redundant. One clear need is to coordinate distributed information gathering activities. For example, if an agent is intending to travel near a site that the surveillance agent intends to observe, it may do just as well to request that the agent (human or robot) make the observation instead. Similarly, if an agent has already observed a site and returned the necessary information to the user, the surveillance agent should be made aware of it. A second, related, facet of this integration/coordination problem is to track the changing

goals and interests of the user and then automatically adapt surveillance strategy in support. For example, the agent may be considering a plan whose outcome depends on knowing the likelihood of an event occurring at a given location. The vehicle could support the planning process by including that location in its surveillance schedule in order to build up a statistical profile of the area.

Coordination cannot be achieved simply by providing a better interface to the vehicle. It requires constructing technologies and procedures that support a shared environment, conceptually centralized even if physically distributed, for commanding and controlling resources, planning system wide activities and communicating activities and intentions to other agents. The need for this kind of environment is well-recognized in certain communities, particularly in the armed forces where highly integrated processes for command, control and communication have long been a key concept of operations. More generally, this kind of integration represents a natural progression for many technologies including surveillance and reconnaissance agents. As UAVs go from costly, rare and highly specialized assets to relatively inexpensive commodities used for long periods and diverse purposes, infrastructure will be developed to facilitate their integration into the larger operational environment.

Future Work

Our focus so far has been on surveillance decision-making and on characterizing the best role for humans in this process. We expect to extend this work to consider more diverse factors and thus refine our understanding of how to allocate decision-making responsibilities between humans and machines. As an extension to this effort, we expect to design decision aids to help human operators make surveillance decisions in operational contexts where no effective algorithmic approach is available. The evaluation framework developed to compare algorithms to one another and to human operator performance should apply equally to the evaluation of human operators with and without candidate decision aids.

We are at a very early stage in our thinking about how to address the other human interaction problems identified above. A range of techniques have been developed to help elicit utility knowledge from human experts (French 1986; Wang and Boutilier 2003), possibly providing a way to address the problem of acquiring domain models from users. Dialogue management technologies have been applied successfully to the control of complex systems, and may prove effective for integrating user-initiated reconnaissance activities with ongoing surveillance. Numerous new and emerging technologies are likely to prove invaluable for coordinating the activities of surveillance vehicle with other system activities. For example, collaborative work environments

might provide a framework for passively inferring user intentions and candidate plans, allowing a system to unobtrusively direct UAV reconnaissance and surveillance effort to the sites of greatest current relevance to the user. In general, UAV-based autonomous surveillance presents a range of open problems for diverse communities concerned with human-machine interaction.

Acknowledgements

This work was supported by the NASA CICT / Intelligent Systems program. We thank Michael Shafto and Jeremy Frank for many useful discussions.

References

- DoD (2002) *Unmanned aerial vehicles roadmap 2002-2027*. Report to the Office of the Secretary of Defense.
- Freed, M. (1998) Managing multiple tasks in complex, dynamic environments. In *Proceedings of the 1998 National Conference on Artificial Intelligence*. Madison, Wisconsin.
- Freed, M. (2000) Reactive Prioritization. *Proceedings 2nd NASA International Workshop on Planning and Scheduling for Space*. San Francisco, CA.
- Freed, M., Matessa, M., Remington, R. and Vera, A. (2003) How Apex Automates CPM-GOMS. In *Proceedings of the 2003 International Conference on Cognitive Modeling*. Bamberg, Germany.
- Freed, M., Harris, R. and Shafto, M.G. (2004) Comparing methods for UAV-based autonomous surveillance. Submitted for publication.
- French, S. (1986) *Decision Theory*. Halsted Press, New York.
- John, B., Vera, A., Matessa, M., Freed, M. and Remington, R. (2002) Automating CPM-GOMS. In *CHI 2002 Conference Proceedings*.
- Massios, N., Dorst, L. and Voorbraak, F. A Strategy for Robot Surveillance Using the Hierarchical Structure of the Environment. In *Proceedings of the IJCAI 2001 Workshop on Reasoning with Uncertainty in Robotics*, 43-50, Seattle, USA.
- Reinelt, G. (1994) *The Travelling Salesman. Computational Solutions for TSP Applications*, volume 840 of *Lecture Notes in Computer Science*. Springer-Verlag, Berlin Heidelberg New York, 1994.
- Wang, T. and Boutilier, C. (2003) Incremental utility elicitation with the minimax regret decision criterion. In *Proceedings of the 2003 International Joint Conference on Artificial Intelligence*. Acapulco, Mexico.
- Whalley, M., Freed, M., Takahashi, M., Christian, D., Patterson-Hine, A., Schulein, G. and Harris R. (2003) The NASA/Army Autonomous Rotorcraft Project. In *Proceedings of the American Helicopter Society 59th Annual Forum*. Phoenix, Arizona.